

Formation of Robust Multi-Agent Networks Through Self-Organizing Random Regular Graphs

A. Yasin Yazıcıoğlu, Magnus Egerstedt, and Jeff S. Shamma

Abstract—Multi-agent networks are often modeled as interaction graphs, where the nodes represent the agents and the edges denote some direct interactions. The robustness of a multi-agent network to perturbations such as failures, noise, or malicious attacks largely depends on the corresponding graph. In many applications, networks are desired to have well-connected interaction graphs with relatively small number of links. One family of such graphs is the random regular graphs. In this paper, we present a decentralized scheme for transforming any connected interaction graph with a possibly non-integer average degree of k into a connected random m -regular graph for some $m \in [k, k + 2]$. Accordingly, the agents improve the robustness of the network with a minimal change in the overall sparsity by optimizing the graph connectivity through the proposed local operations.

I. INTRODUCTION

Multi-agent networks have been used to characterize a large number of natural and engineered systems such as biological systems, financial networks, social networks, communication systems, transportation systems, power grids, and robotic swarms. Multi-agent networks are often represented via their interaction graphs, where the nodes correspond to the agents and the edges exist between the agents having some direct interaction. Various system properties such as the robustness, mixing time, or controllability of a network greatly depend on the structure of the corresponding graph (e.g., [1], [2], [3], [4], [5]). Therefore, graph theoretic analysis of networked systems has received a considerable amount of attention during the last decade (e.g., [6], [7]).

In many applications, multi-agent networks face various perturbations such as component failures, noise, or malicious attacks. One of the fundamental measures that capture the robustness of networks to such perturbations is the connectivity of the corresponding graph. A graph is said to be k -node (or k -edge) connected if at least k nodes (or edges) should be removed to render the graph disconnected. In general, graphs with higher connectivity have higher robustness to the targeted removal of their components [1], [2]. For many networks, well-connected interaction graphs provide robustness to not only the structural failures but also the functional

perturbations such as noise (e.g., [8], [9]). Connectivity can also be quantified via some spectral measures such as the algebraic connectivity [10] or the Kirchhoff index [11]. An arguably richer measure of connectivity is the edge (or node) expansion ratio (e.g., [12], [13]). If the expansion ratio of a graph is small, then it is possible to disconnect a large set of nodes by removing only a small number of edges (or nodes). Graphs with high expansion ratios are called expanders. A detailed overview of expanders and their numerous applications can be found in [14] and the references therein.

The connectivity of a network can be improved by adding more edges to the graph. However, each edge stands for some communications, sensing, or a physical link between the corresponding agents. Hence, more edges require more resources. Moreover, too many edges may lead to higher vulnerability to the cascading failures such as epidemics (e.g., [15], [16]). Consequently, having a small number of edges (i.e., sparsity) is desired in many applications. One family of well-connected yet sparse graphs is the random regular graphs. A graph is called an m -regular graph if the number of edges incident to each node (the degree) is equal to m . A random m -regular graph of order n is a graph that is selected uniformly at random from the set of all m -regular graphs with n nodes. As n goes to infinity, almost every m -regular graph is an expander for any $m \geq 3$ [13], [17].

In this paper, we present a decentralized scheme to transform any connected interaction graph into a connected random regular graph with a similar number of edges as the initial graph. Some preliminary results of this work were presented in [18] and [19], where any initial graph with an integer average degree, $k \in \mathbb{N}$, was transformed into a random k -regular graph. In this paper, we extend our earlier works to obtain a method that is applicable to the most generic case, i.e. $k \in \mathbb{R}$. The proposed method transforms any graph with a possibly non-integer average degree of k into a random m -regular graph for some $m \in [k, k + 2]$. As such, the graph becomes well-connected, regardless of the network size.

The organization of this paper is as follows: Section II presents some related work. Section III provides some graph theory preliminaries. Section IV formulates the problem and presents the proposed solution. Section V provides a distributed implementation and an analysis of the resulting dynamics. Section VI provides some simulation results to demonstrate the performance of the proposed solution. Finally, Section VII concludes the paper.

This work was supported by AFOSR/MURI #FA9550-10-1-0573 and ONR Project #N00014-09-1-0751.

A. Yasin Yazıcıoğlu is with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, yasiny@mit.edu.

Magnus Egerstedt is with the School of Electrical and Computer Engineering, Georgia Institute of Technology, magnus@gatech.edu.

Jeff S. Shamma is with the School of Electrical and Computer Engineering, Georgia Institute of Technology, shamma@gatech.edu, and with King Abdullah University of Science and Technology (KAUST), jeff.shamma@kaust.edu.sa.

II. RELATED WORK

In some applications, the robustness of a graph is related to the centrality measures such as the degree, betweenness, closeness, and eigenvector centralities. Loosely speaking, the centrality measures capture the relative importances of the nodes in a graph. Detailed reviews on the centrality measures and their applications can be found in [20], [21] and the references therein. Typically, the perturbations applied to the nodes with higher centrality scores have a stronger impact on the overall system (e.g., [22], [23], [24], [25]). Hence, graphs with unbalanced centrality distributions are usually vulnerable to such worst-case perturbations.

In the literature, there are many works related to the construction of robust interaction graphs. Some of these works consider how a robustness measure can be improved via some modifications to the graph topology. For instance, such improvement can be achieved by rewiring a small percentage of the existing edges (e.g., [26]) or adding a small number of edges to the graph (e.g., [27], [9]).

Another group of studies consider the explicit construction of expanders. Expanders can be constructed via graph operations such as zig-zag product (e.g., [28], [29]), or derandomized graph squaring [30]. Furthermore, for any $m \in \mathbb{N}$ such that $m-1$ is a prime power, an explicit algebraic construction method for a family of m -regular expanders, i.e. Ramanujan graphs, was presented in [31]. In [32], Watts-Strogatz small-world networks are transformed into quasi Ramanujan graphs by rewiring some of the edges.

Expanders can also be built as random m -regular graphs. A detailed survey of the various models of random regular graphs as well as their properties can be found in [33] and the references therein. As n goes to infinity, almost every m -regular graph has an algebraic connectivity arbitrarily close to $m - 2\sqrt{m-1}$ for $m \geq 3$ [13], [17]. For regular graphs, high algebraic connectivity implies high expansion ratios (e.g., [34]). Hence, for any fixed $m \geq 3$, a random m -regular graph has the algebraic connectivity and expansion ratios bounded away from zero, even if the network size is arbitrarily large.

A random m -regular graph with n nodes can be constructed by generating m copies for each node, picking a uniform random perfect matching on the nm copies, and connecting any two nodes if the matching contains an edge between their copies (e.g., [35], [36]). In [37], the authors present a distributed scheme for incrementally building random $2m$ -regular multi-graphs with m Hamiltonian cycles. Alternatively, some graph processes may be designed to transform an initial m -regular graph into a random m -regular graph by inducing a Markov chain with a uniform limiting distribution over the set of m -regular graphs (e.g., [38], [39]). The method in this paper is also based on designing a graph process with a uniform limiting distribution over the set of m -regular graphs. Compared to the similar works in the literature, the proposed scheme is applicable to the most generic case, and it is decentralized. The initial graph is not required to satisfy some strong properties such as being

regular or having an integer average degree. Furthermore, the global transformation is achieved via only some local graph operations.

III. PRELIMINARIES

An undirected graph, $\mathcal{G} = (V, E)$, consists of a set of nodes, V , and a set of edges, E , given by unordered pairs of nodes. A graph is connected if there exists a path between any pair of nodes. A path is a sequence of nodes such that an edge exists between any two consecutive nodes in the sequence. Any two nodes are said to be adjacent if an edge exists between them. We refer to the set of nodes adjacent to any node, $i \in V$, as its neighborhood, \mathcal{N}_i , defined as

$$\mathcal{N}_i = \{j \mid (i, j) \in E\}. \quad (1)$$

For any node i , the number of nodes in its neighborhood is called its degree, d_i , i.e.,

$$d_i = |\mathcal{N}_i|, \quad (2)$$

where $|\mathcal{N}_i|$ denotes the cardinality of \mathcal{N}_i . For any graph \mathcal{G} , we use $\delta(\mathcal{G})$, $\Delta(\mathcal{G})$ and $\bar{d}(\mathcal{G})$ to denote the minimum, the maximum, and the average degrees, respectively. We refer to the difference of the maximum and the minimum node degrees in a graph as the degree range, $f(\mathcal{G})$, i.e.

$$f(\mathcal{G}) = \Delta(\mathcal{G}) - \delta(\mathcal{G}). \quad (3)$$

For any undirected graph, $\mathcal{G} = (V, E)$, the graph Laplacian, L , is a symmetric matrix whose entries are given as

$$[L]_{ij} = \begin{cases} d_i & \text{if } i = j, \\ -1 & \text{if } (i, j) \in E, \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

The second-smallest eigenvalue of the graph Laplacian is known as the algebraic connectivity of the graph, $\alpha(\mathcal{G})$.

Local graph transformations can be represented using the framework of graph grammars (e.g., [40]). A grammar, Φ , is a set of rules, where each rule is defined as a label-dependent graph transformation. Each rule is represented as an ordered pair of labeled graphs, $r = (g_l, g_r)$, where the labels represent the node states. Graph grammars operate on labeled graphs. A labeled graph, $\mathcal{G} = (V, E, l)$, consists of a node set, V , an edge set, E , and a labeling function, $l : V \mapsto \Sigma$, where Σ is the set of feasible node labels. A rule is said to be applicable to a labeled graph, $\mathcal{G} = (V, E, l)$, if \mathcal{G} has a subgraph isomorphic to g_l , i.e. if there is a bijection, which preserves node labels and edges, between g_l and a subgraph of \mathcal{G} . A rule, $r = (g_l, g_r)$, transforms any graph isomorphic to g_l to a graph isomorphic to g_r . A labeled initial graph, $\mathcal{G}(0)$, along with a grammar, Φ , defines a non-deterministic system represented as $(\mathcal{G}(0), \Phi)$.

IV. DECENTRALIZED FORMATION OF RANDOM REGULAR GRAPHS

A. Problem Formulation

Motivated by the connectivity properties of almost every m -regular graph for $m \geq 3$, this paper is focused on the

following problem: Assume that a multi-agent network is initialized with an arbitrary connected interaction graph. How can the agents reconfigure their links locally in a decentralized fashion such that the resulting dynamics transform the interaction graph into a connected random regular graph with a similar sparsity as the initial graph? Such a transformation is illustrated in Fig. 1.




Fig. 1. A fragile interaction graph (a) with an average degree 2.8 is transformed into a random 3-regular graph (b). The graphs have similar sparsities, whereas the structure in (b) has a significantly improved connectivity.


In order to transform the interaction graph into a random regular graph, the agents need to pursue three global objectives: 1) balance the degree distribution, 2) randomize the links, and 3) drive the average degree to an integer close to its initial value. Furthermore, the agents should ensure that the graph remains connected as they modify their links. This paper presents a set of locally applicable graph reconfiguration rules to pursue all of these global objectives in parallel. In the remainder of this section, we will incrementally build the proposed scheme using the framework of graph grammars.

B. Degree Balancing

As the first step towards building random regular graphs, we present a local rule for balancing the degree distribution in a network while maintaining the connectivity and the total number of edges. The degree balancing task can be considered as a quantized consensus problem (e.g., [41], [42]), where each local update needs to be realized via some local changes in the structure of the graph. In this part, we design a single-rule grammar, Φ_R , for balancing the degree distribution. In Φ_R , each node is labeled with its degree, i.e.

$$l(i) = d_i, \quad \forall i \in V. \quad (5)$$

The proposed grammar, Φ_R , is defined as



where d_i , d_j , and d_h denote the degrees of the corresponding nodes.

In accordance with Φ_R , nodes behave as follows: Let i and j be any pair of adjacent nodes. If $d_j < d_i$, then a new link is formed between j and an arbitrary neighbor of i , say h , that is not currently linked with j . At the same time, the link between i and h is terminated to maintain the overall sparsity. Note that if $d_j < d_i$, then i has at least one such exclusive neighbor, h . Furthermore, if a connected graph is not regular then there always exists at least one pair of adjacent nodes, i and j , such that $d_i > d_j$. Hence, any connected graph is stationary (i.e., Φ_R is not applicable anywhere on the graph) if and only if it is a regular graph. Moreover, any concurrent application of Φ_R on a connected graph, \mathcal{G} , maintains the connectivity and the average degree of \mathcal{G} .

A detailed analysis of the dynamics induced by Φ_R can be found in [18], where it is shown that the degree range, $f(\mathcal{G})$, monotonically decreases and converges to its minimum feasible value under Φ_R . As such, the graph converges to a regular graph (i.e., $f(\mathcal{G}) = 0$) if the average degree is an integer. Otherwise, $f(\mathcal{G})$ converges to 1. For the sake of brevity, we skip the details of this analysis and only provide its main result in Theorem 4.1.

Theorem 4.1 *Let $\mathcal{G}(0)$ be a connected graph and let $\tau = \{\mathcal{G}(0), \mathcal{G}(1), \dots\}$ be a feasible trajectory of $(\mathcal{G}(0), \Phi_R)$. Then, $\tau_f = \{f(\mathcal{G}(0)), f(\mathcal{G}(1)), \dots\}$, almost surely converges to an integer, τ_f^* , such that*

$$\tau_f^* = \begin{cases} 0 & \text{if } \bar{d}(\mathcal{G}(0)) \in \mathbb{N} \\ 1 & \text{otherwise,} \end{cases} \quad (6)$$

where $\bar{d}(\mathcal{G}(0))$ denotes the average degree of $\mathcal{G}(0)$.

The grammar Φ_R transforms any connected graph with an integer average degree, $k \in \mathbb{N}$, into a connected k -regular graph. Note that, although almost every k -regular graph is an expander for $k \geq 3$, Φ_R may still result in a rare configuration with an arbitrarily small expansion rate. For instance, both of the 3-regular graphs in Fig. 2 are stationary under Φ_R whereas the graph in Fig. 2a can have half of the network disconnected due to the failure of a single edge or a single node. The probability of converging to an undesired equilibrium such as the graph in Fig. 2a depends on the initial graph. In the next part, we will build on Φ_R to avoid such undesired outcomes.





Fig. 2. A poorly-connected 3-regular graph (a) and a well-connected 3-regular graph (b).

C. Link Randomization

In order to ensure that the interaction graph does not converge to a poorly-connected configuration, we extend Φ_R by adding a second rule for randomizing the links between the agents. The purpose of this additional rule is to induce a uniform limiting distribution over the set of all connected k -regular graphs instead of having the network converge to an element of that set. As such, the resulting graph will be a random k -regular graph. A locally applicable method was introduced in [39] for transforming regular graphs into random regular graphs. We define an additional rule for Φ_R based on this method. The resulting grammar, Φ_{RR} , uses the same labels as Φ_R , and it is defined as



where d_i , d_j , d_h , and d_f denote the degrees of the corresponding nodes. In accordance with Φ_{RR} , if a node has more links than one of its neighbors has, then it rewires one of its other neighbors to its less-connected neighbor (r_1). Also, adjacent nodes exchange their exclusive neighbors (r_2).

Both rules in Φ_{RR} maintain the number of edges and the connectivity. Unlike Φ_R , a connected graph is never stationary under Φ_{RR} (unless it is a complete graph). Once a regular graph is reached, r_1 is not applicable anymore, but the agents keep randomizing their links via r_2 .

There may be many feasible concurrent applications of Φ_{RR} on an interaction graph. In such cases, the agents need to execute one of the alternatives randomly. These randomizations do not have to be uniform, and assigning each alternative a non-zero probability is sufficient for the desired limiting behavior to emerge. We provide Algorithm I as a such distributed implementation that leads to a uniform limiting distribution over the set of connected k -regular graphs.

Algorithm I is memoryless since each iteration only depends on the current graph, and the probability of any feasible transition is independent of the past. As such, it induces a Markov chain, P_{RR} , over the set of connected graphs with the same average degree as the initial graph, i.e.

$$\mathbb{G}_{n,k} = \{\mathcal{G} = (V, E) \mid |V| = n, |E| = \frac{kn}{2}\}, \quad (7)$$

where $k = \bar{d}(\mathcal{G}(0))$. Let $\mathbb{G}_{n,k}^0 \subseteq \mathbb{G}_{n,k}$ be the set of connected k -regular graphs. Note that if $k \in \mathbb{N}$, then $\mathbb{G}_{n,k}^0 \neq \emptyset$. A thorough analysis of the dynamics induced by Algorithm I

along with the following result can be found in [19].

Theorem 4.2 *Let $\mathbb{G}_{n,k}$ satisfy $k \in \mathbb{N}$. Then, P_{RR} has a unique limiting distribution, μ^* , satisfying*

$$\mu^*(\mathcal{G}) = \begin{cases} 1/|\mathbb{G}_{n,k}^0| & \text{if } \mathcal{G} \in \mathbb{G}_{n,k}^0, \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

Algorithm I: Distributed Implementation of Φ_{RR}

```

1 : initialize:  $\mathcal{G} = (V, E)$  connected,  $\epsilon \in (0, 1)$ 
2 : repeat
3 :   Each agent,  $i$ , is active with probability  $1 - \epsilon$ .
4 :   Each active  $i$  picks a random  $j \in \mathcal{N}_i$ .
5 :   For each  $i$ ,  $R_i = \{i' \in \mathcal{N}_i \mid i' \text{ picked } i\}$ .
6 :   for (each  $(i, j)$  s.t.  $i \in R_j, j \in R_i, d_i \geq d_j$ )
7 :      $\max\{i, j\}$  picks a random  $r \in \Phi_{RR}$ .
8 :     if ( $r = r_1, d_i > d_j, |R_i| \geq 2$ )
9 :        $i$  picks a random  $h \in R_i \setminus \{j\}$ .
10 :      if  $((j, h) \notin E)$ 
11 :         $E = (E \setminus \{(i, h)\}) \cup \{(j, h)\}$ .
12 :      end if
13 :     else if ( $r = r_2, |R_i| \geq 2, |R_j| \geq 2$ )
14 :        $i$  picks a random  $h \in R_i \setminus \{j\}$ .
15 :        $j$  picks a random  $f \in R_j \setminus \{i\}$ .
16 :       if  $((i, f) \notin E, (j, h) \notin E)$ 
17 :         $E = (E \setminus \{(i, h), (j, f)\}) \cup \{(i, f), (j, h)\}$ .
18 :       end if
19 :     end if
20 :   end for
21 : end repeat

```

In light of Theorem 4.2, any connected initial graph with an integer average degree, $k \in \mathbb{N}$, can be transformed into a connected random k -regular graph via Φ_{RR} . For any $k \geq 3$, such graphs have their algebraic connectivity and expansion ratios bounded away from zero, even when the network size is arbitrarily large. If $k \notin \mathbb{N}$, then the degree range converges to 1 via r_1 , as given in Theorem 4.1. In that case, the limiting graph can be loosely described as a random almost-regular graph. While such graphs have many structural similarities to random regular graphs, there is not an exact characterization of their expansion properties. In the next part, we will build on Φ_{RR} to obtain a grammar that produces random regular graphs, even when $k \notin \mathbb{N}$.

D. Average Degree Manipulation


Both Φ_R and Φ_{RR} were designed to maintain the overall network sparsity. In order to obtain random regular graphs even when the initial graph has a non-integer average degree, $k \notin \mathbb{N}$, the agents need to manipulate the total number of edges. To this end, the sparsity constraint should be relaxed to allow for occasional increases and decreases in the number of edges as long as the graph is not regular and the average degree is within some proximity of its initial value. Also, the feasible range of the number of edges should ensure that the average degree can reach an integer value. Since an m -regular graph with an odd number of nodes does not exist

if m is odd, an even integer should always exist within the feasible range of the average degree. In this part, we derive such a grammar, Φ^* , by building on Φ_{RR} .

In Φ^* , each node, $i \in V$, has a label, $l(i)$, consisting of two elements. The first entry of the label is equal to the degree of the node, just as in Φ_R and Φ_{RR} . The second entry of $l(i)$ is a binary flag that denotes whether the corresponding agent is allowed to locally increase or decrease the number of edges in the network. A node, i , can add an extra edge to the graph only if $w_i = 0$, and it can remove an edge only if $w_i = 1$. As such, the node labels are defined as

$$l(i) = \begin{bmatrix} d_i \\ w_i \end{bmatrix}, \forall i \in V, \quad (9)$$

where d_i denotes the degree of i , and $w_i \in \{0, 1\}$. For any graph, the degree of each node is already encoded through the edge set, E . Hence, for the sake of simplicity, a labeled graph will be denoted as $\mathcal{G} = (V, E, w)$ in the sequel. Using the node labels given in (9), the proposed grammar, $\Phi^* = \{r_1, r_2, r_3, r_4\}$, is defined as



The rules in Φ^* can be interpreted as follows: The first rule, r_1 , is a variant of the degree balancing rule in Φ_{RR} . While r_1 induces a transformation identical to the degree balancing rule in Φ_{RR} with a probability, $1 - \beta$, there is also a nonzero probability, β , that if the higher degree node has its binary flag equal to 0, then it maintains its link with the neighbor introduced to its lower degree neighbor. Note that the total number of edges increases in that case, and the binary flag of the corresponding higher degree node becomes 1. The second rule, r_2 , is the same link randomization rule as in Φ_{RR} . The third rule, r_3 , is the edge removal rule which

breaks one edge of a triangle if all of its nodes do not have the same degree and one of the higher degree nodes has its binary flag equal to 1. Note that the total number of edges decreases in that case, and the binary flag of the corresponding higher degree node becomes 0. The final rule, r_4 , is for exchanging the binary flags among adjacent nodes. The purpose of r_4 is to ensure that the edge additions and removals will be applicable as long as they are needed to drive the average degree to an integer. In the remainder of this section, the dynamics induced by Φ^* is analyzed.

Lemma 4.3 *Graph connectivity is maintained under any concurrent application of Φ^* .*

Proof: Let a connected graph $\mathcal{G}(t)$ be transformed into $\mathcal{G}(t+1)$ via some concurrent application of Φ^* . Since $\mathcal{G}(t)$ is connected, for every pair of nodes, $v, v' \in V$, $\mathcal{G}(t)$ contains a finite path P between v and v' . If all the edges traversed in P are maintained in $\mathcal{G}(t+1)$, then P is also a valid path on $\mathcal{G}(t+1)$. Otherwise, any missing edge is removed due to an execution of r_1 , r_2 , or r_3 . Accordingly, the following procedure can be executed to transform P into a valid path on $\mathcal{G}(t+1)$ between the same terminal nodes as P :

For each edge that is traversed in P but missing in $\mathcal{G}(t+1)$,

1) If the edge was removed via r_1 or r_3 , then replace the corresponding part of the path, $\{i, h\}$, with $\{i, j, h\}$.

2) If the edge was removed via r_2 , then replace the corresponding part of the path, $\{i, h\}$ (or $\{j, f\}$), with $\{i, j, h\}$ (or $\{j, i, f\}$).

Consequently, if $\mathcal{G}(t)$ is connected, then $\mathcal{G}(t+1)$ is also connected. ■

Note that, unlike Φ_R and Φ_{RR} , the number of edges is not maintained under Φ^* due to the possible edge additions and removals. However, the number of edges is bounded in an interval through the binary labels, w .

Lemma 4.4 *Let $\mathcal{G} = (V, E, w)$ be a graph, and let τ be any feasible trajectory of (\mathcal{G}, Φ^*) . For any $\mathcal{G}' = (V, E', w') \in \tau$,*

$$|E'| - |E| = \mathbf{1}^T w' - \mathbf{1}^T w, \quad (10)$$

where $\mathbf{1}$ is the vector of all ones.

Proof:

Let $\mathcal{G}(t) = (V, E(t), w(t))$ and $\mathcal{G}(t+1) = (V, E(t+1), w(t+1))$ be a pair of consecutive graphs on any trajectory, τ , of (\mathcal{G}, Φ^*) . Note that the number of edges is maintained under r_2 and r_4 , and it can increase by 1 only due to r_1 or decrease by 1 only due to r_3 . If an application of r_1 adds an extra edge to the network, then one of the nodes involved in that transformation, i , satisfies $w_i(t+1) - w_i(t) = 1$, whereas the other two participating nodes maintain their w entries. On the other hand, for each application of r_3 , one of the nodes involved in that transformation, i , satisfies $w_i(t+1) - w_i(t) = -1$ while the other two participating nodes maintain their w entries. Furthermore, any rule application that maintains the number of edges also maintains the sum of w entries of the

participating nodes. In particular, any edge rewiring via r_1 or any neighbor swapping via r_2 maintain the w entries of the participating nodes, whereas any label exchange via r_4 maintains the sum of labels. Hence, $E(t)$ and $E(t+1)$ satisfy

$$|E(t+1)| - |E(t)| = \mathbf{1}^T w(t+1) - \mathbf{1}^T w(t). \quad (11)$$

Using induction, it can be shown that (11) implies (10) for any $\mathcal{G}' = (V, E', w') \in \tau$. ■

Corollary 4.5 Let $\mathcal{G} = (V, E, \mathbf{0})$ be a graph, and let τ be any feasible trajectory of (\mathcal{G}, Φ^*) . For any $\mathcal{G}' = (V, E', w') \in \tau$,

$$|E| \leq |E'| \leq |E| + |V|. \quad (12)$$

Proof: Let $\mathcal{G} = (V, E, \mathbf{0})$ be a graph, and let τ be any feasible trajectory of (\mathcal{G}, Φ^*) . In light of Lemma 4.4, every $\mathcal{G}' = (V, E', w') \in \tau$ satisfies

$$|E'| - |E| = \mathbf{1}^T w' - \mathbf{1}^T \mathbf{0} = \mathbf{1}^T w'. \quad (13)$$

Since each entry of w is either equal to 1 or equal 0,

$$0 \leq \mathbf{1}^T w' \leq |V|. \quad (14)$$

Hence, (13) and (14) together imply (12). ■

Corollary 4.5 provides the upper and the lower bounds on the of number of edges along any trajectory of (\mathcal{G}, Φ^*) for any $\mathcal{G} = (V, E, \mathbf{0})$. Lemma 4.6 shows that the corresponding interval always contains a value that implies an integer average degree.

Lemma 4.6 For any $\mathcal{G} = (V, E)$, a graph, $\mathcal{G}' = (V, E')$, satisfying $\bar{d}(\mathcal{G}') \in \mathbb{N}$ can be formed by adding at most $|V|$ edges to \mathcal{G} .

Proof: Let $\mathcal{G} = (V, E)$ be a graph, and let $\mathcal{G}' = (V, E')$ be a graph that is formed by adding η edges to \mathcal{G} . Then,

$$\bar{d}(\mathcal{G}') = \frac{2(|E| + \eta)}{|V|}. \quad (15)$$

For any $\mathcal{G} = (V, E)$, there exists unique pair of integers, $p, q \in \mathbb{N}$, such that $0 \leq q < |V|$ and

$$|E| = p|V| + q. \quad (16)$$

Let $\eta^* = |V| - q$. Note that $\eta^* \leq |V|$ since $0 \leq q < |V|$. Furthermore, by plugging η^* into (15), we obtain

$$\frac{2(|E| + \eta^*)}{|V|} = 2p + 2\frac{q}{|V|} + 2 - 2\frac{q}{|V|} = 2(p+1). \quad (17)$$

Since $p \in \mathbb{N}$, we have $2(p+1) \in \mathbb{N}$. Hence, for any $\mathcal{G} = (V, E)$, there exists $0 \leq \eta \leq \eta^*$ such that a graph, $\mathcal{G}' = (V, E')$, satisfying $\bar{d}(\mathcal{G}') \in \mathbb{N}$ can be obtained by adding η edges to \mathcal{G} . ■

In light of Lemma 4.6, the interval in (12) always has at least one value implying an integer average degree. Next, it is shown that (\mathcal{G}, Φ^*) almost surely reaches a regular graph for any connected $\mathcal{G} = (V, E, \mathbf{0})$ such that $\bar{d}(\mathcal{G}) > 2$.

Lemma 4.7 Let $\mathcal{G} = (V, E, w)$ be a connected graph such that $\bar{d}(\mathcal{G}) \notin \mathbb{N}$. If $\mathbf{1}^T w < |V|$, then \mathcal{G} can be transformed via Φ^* into a graph, $\mathcal{G}' = (V, E', w')$, satisfying $|E'| = |E| + 1$.

Proof: If $\bar{d}(\mathcal{G}) \notin \mathbb{N}$, then \mathcal{G} has to be a non-regular graph. Furthermore, since \mathcal{G} is connected, then there has to be at least a pair of neighboring nodes, i and j , with different degrees. Without loss of generality let $d_i > d_j$. Then, i should have at least one neighbor that is not linked to j . As such, r_1 is applicable on $\mathcal{G} = (V, E)$. Furthermore, since each entry of w is either 0 or 1, at least one entry of w is equal to 0 if $\mathbf{1}^T w < |V|$. If $w_i = 0$, then an extra edge can be added to the graph via r_1 with probability β . If $w_i \neq 0$, then there exists another node, i' , such that $w_{i'} = 0$. Apply r_4 along the shortest path between i and i' to obtain $w_i = 0$. Once $w_i = 0$ is obtained, then with probability β an extra edge can be added to the graph in accordance with r_1 to obtain a graph, $\mathcal{G}' = (V, E', w')$, satisfying $|E'| = |E| + 1$. ■

Lemma 4.8 Let $\mathcal{G} = (V, E, w)$ be a connected triangle-free graph. If $\bar{d}(\mathcal{G}) > 2$, then \mathcal{G} can be transformed via Φ^* into a graph, $\mathcal{G}' = (V, E', w)$, containing at least one triangle.

Proof: Let $\mathcal{G} = (V, E, w)$ be a connected triangle-free graph. Since $\bar{d}(\mathcal{G}) > 2$, \mathcal{G} has to be a cyclic graph. Consider the shortest cycle on \mathcal{G} , C_s . Note that C_s has to contain at least 4 nodes, since \mathcal{G} is triangle-free. Furthermore, if $\bar{d}(\mathcal{G}) > 2$, then C_s cannot contain all the nodes of \mathcal{G} , as otherwise \mathcal{G} has to be a cycle graph with $\bar{d}(\mathcal{G}) = 2$. Also, since \mathcal{G} is a connected graph, at least one node on C_s must be linked to a node that is not contained in C_s . Without loss of generality, let i be a node on C_s connected to an off-cycle node, h . Furthermore, let j denote a neighbor of i on C_s . Note that h can not be connected to j since \mathcal{G} is triangle-free. Furthermore, since the C_s contains at least 4 nodes, j has to have another neighbor on C_s , f , which is not connected to i . If i and j execute r_2 by rewiring h and f to each other, C_s becomes a shorter cycle. Note that w is invariant under r_2 . Hence, by repeating this process until C_s consists of 3 nodes, \mathcal{G} can be transformed into a graph, $\mathcal{G}' = (V, E', w)$, containing at least one triangle. ■

Lemma 4.9 Let $\mathcal{G} = (V, E, w)$ be a connected graph such that $\bar{d}(\mathcal{G}) \notin \mathbb{N}$. If $\mathbf{1}^T w > 0$ and $\bar{d}(\mathcal{G}) > 3$, then \mathcal{G} can be transformed via Φ^* into a graph, $\mathcal{G}' = (V, E', w')$, satisfying $|E'| = |E| - 1$.

Proof: Let $\mathcal{G} = (V, E, w)$ be a connected graph such that $\bar{d}(\mathcal{G}) \notin \mathbb{N}$. Since $\bar{d}(\mathcal{G}) \notin \mathbb{N}$, the degree range can be reduced to 1 without changing \bar{d} and w via r_1 , as given in Theorem 4.1. Furthermore, in light of Lemma 4.8, if the resulting graph is triangle-free, then it can be transformed via r_2 into a graph, $\mathcal{G}^+ = (V, E^+, w)$, which contains at least one triangle. Since $\Delta(\mathcal{G}^+) - \delta(\mathcal{G}^+) = 1$ and $\bar{d}(\mathcal{G}^+) = \bar{d}(\mathcal{G}) > 3$, we have $\delta(\mathcal{G}^+) \geq 3$.

If $\mathbf{1}^T w > 0$, then at least one entry of w is equal to 1. As such, $w_i = 1$ can be reached for any $i \in V$ via a

sequence of r_4 applications. Hence, if any of the triangles in \mathcal{G}^+ involves nodes with non-uniform degrees, then one edge of the triangle can be broken in accordance with r_3 to obtain a graph, $\mathcal{G}' = (V, E')$, satisfying $|E'| = |E| - 1$.

On the other hand, assume that all the triangles on \mathcal{G}^+ consist of nodes with uniform degrees. Let i, j , and h be the nodes of such a triangle. Consider the shortest path from any node of this triangle to a node whose degree is not equal to the degrees of the triangle nodes. Without loss of generality, let this path be between i and i' such that $d_{i'}^+ \neq d_i^+$. Note that each node other than i' on this path has its degree equal to d_i^+ , and either $d_i^+ > d_{i'}^+$ or $d_i^+ < d_{i'}^+$. Both cases are inspected below:

1) If $d_i^+ > d_{i'}^+$, then $d_i^+ \geq 4$ given that $\delta(\mathcal{G}^+) \geq 3$. Hence, the node next to i' on the shortest path has at least one neighbor outside the shortest path that is not connected to i' . Applying r_1 by rewiring such a neighbor to i' results in a shorter path from i to a smaller degree node. Applying the same procedure eventually results in a smaller degree node being adjacent to i . Since $d_i^+ \geq 4$, i has at least one neighbor other than j and h to rewire to its smaller degree neighbor. After the corresponding application of r_1 , we obtain a triangle of nodes with non-uniform degrees. Hence, one edge of the triangle can be broken in accordance with r_3 to obtain a graph, $\mathcal{G}' = (V, E', w')$, satisfying $|E'| = |E| - 1$.

2) If $d_i^+ < d_{i'}^+$, then i' can apply r_1 to rewire one of its neighbors to the node on the shortest path next to itself. This results in a shorter path from i to higher degree node is obtained. Applying the same procedure eventually increases the degree of i , and we obtain a triangle of nonuniform degrees. Hence, one edge of the triangle can be broken in accordance with r_3 to obtain a graph, $\mathcal{G}' = (V, E', w')$, satisfying $|E'| = |E| - 1$. ■

Lemma 4.10 *Let $\mathcal{G} = (V, E, 0)$ be a connected graph. If $\bar{d}(\mathcal{G}) > 2$, then (\mathcal{G}, Φ^*) almost surely reaches an m -regular graph such that $\bar{d}(\mathcal{G}) \leq m \leq \bar{d}(\mathcal{G}) + 2$.*

Proof: Let $\mathcal{G} = (V, E, 0)$ be a connected graph such that $\bar{d}(\mathcal{G}) > 2$. Due to Corollary 4.5, the number of edges stays in $\{|E|, |E| + 1, \dots, |E| + |V|\}$ along any trajectory of (\mathcal{G}, Φ^*) . In light of Lemma 4.6, this interval contains at least one value, $|E| + \eta$, such that the corresponding average degree, m , is an integer satisfying

$$\bar{d}(\mathcal{G}) \leq m = \frac{2(|E| + \eta)}{|V|} \leq \bar{d}(\mathcal{G}) + 2. \quad (18)$$

As such, $m \geq 3$ since $\bar{d}(\mathcal{G}) > 2$. Let $\mathcal{G}' = (V, E', w')$ be a graph reached from \mathcal{G} via Φ^* such that $\bar{d}(\mathcal{G}') \notin \mathbb{N}$. Then, either $|E'| < |E| + \eta$ or $|E'| > |E| + \eta$. Both cases are inspected below:

1) If $|E'| < |E| + \eta$, then $\mathbf{1}^T w' = |E'| - |E| < |V|$. As such, in light of Lemma 4.7, \mathcal{G}' can have its number of edges increased by 1 via Φ^* . This process can be repeated unless the average degree reaches an integer value.

2) If $|E'| > |E| + \eta$, then $\mathbf{1}^T w' = |E'| - |E| > 0$. Furthermore, $\bar{d}(\mathcal{G}') > 3$ since $m \geq 3$. As such, in light of Lemma 4.9, \mathcal{G}' can have its number of edges decreased by 1 via Φ^* . This process can be repeated unless the average degree reaches an integer value.

Hence, in either case it is possible to add or remove edges to the network via Φ^* until the number of edges implies an integer average degree, $m \in [\bar{d}(\mathcal{G}), \bar{d}(\mathcal{G}) + 2]$. Once an average degree of m is obtained, then the graph can be driven to a m -regular configuration through a sequence of r_1 applications. As such, any \mathcal{G}' has a non-zero probability of reaching an m -regular graph after a finite sequence of Φ^* applications. Consequently, (\mathcal{G}, Φ^*) almost surely reaches an m -regular graph such that $\bar{d}(\mathcal{G}) \leq m \leq \bar{d}(\mathcal{G}) + 2$. ■

Unlike Φ_R and Φ_{RR} , the degree range does not monotonically decrease under Φ^* due to the possible edge addition and removals. For instance, increasing the total number of edges via r_1 would increase the degree range if $d_h = \Delta(\mathcal{G})$ and $\delta(\mathcal{G}) < d_j$. However, if a regular graph is reached, then the graph remains regular since neither r_1 nor r_3 are applicable on a regular graph.

Lemma 4.11 *Let \mathcal{G} be an m -regular graph. Any feasible trajectory of (\mathcal{G}, Φ^*) consists of only m -regular graphs.*

Proof: On an m -regular graph, r_1 and r_3 are not applicable since $d_i = d_j$ for any pair of nodes, i, j . Furthermore, since the node degrees are invariant to the applications of r_2 and r_4 , any feasible trajectory of (\mathcal{G}, Φ^*) consists of only m -regular graphs. ■

Once an m -regular graph is reached, the graph evolves only under r_2 and r_4 . As such, it keeps randomizing within the corresponding set of m -regular graphs. The induced limiting distribution depends on how the agents implement Φ^* . In the next section, we present a distributed implementation that leads to a uniform limiting distribution over the corresponding set of m -regular graphs.

V. DISTRIBUTED IMPLEMENTATION

In this section, we present Algorithm II as a distributed implementation of Φ^* and analyze the resulting dynamics.

In accordance with Algorithm II, the nodes behave as follows: At each iteration, each node is inactive with a probability, $\epsilon \in (0, 1)$. The inactivation probability, ϵ , ensures that any feasible application of Φ^* can be realized through Algorithm II, as it will be shown in Lemma 5.1. Inactive nodes do not participate in any rule execution in that time step. Each active agent, i , picks one of its neighbors, $j \in \mathcal{N}_i$, uniformly at random, and it communicates its degree to that neighbor. Through these communications, each active agent, i , obtains the list of neighbors that picked itself, $R_i \subseteq \mathcal{N}_i$, and checks if it is matched, i.e. if $j \in R_i$. If that is not the case, then i is a follower in that time step, i.e. it will not initiate a rule execution but it will participate if j wants to rewire i to some other node. If $j \in R_i$, then i and j are matched. Each matched pair randomly pick a candidate rule, $r \in \Phi^*$, that they will potentially execute. In Algorithm II,

without loss of generality, the candidate rule is picked by the agent with the larger node ID, i.e. $\max\{i, j\}$. Depending on the chosen rule, $r \in \Phi^*$, the matched pair of nodes, i and j , do one of the following:

- 1) $r = r_1$: If $d_i > d_j$ and $|R_i| \geq 2$, then i chooses a neighbor, $h \neq j \in R_i$, uniformly at random. If $h \notin \mathcal{N}_j$, then, with probability $1 - \beta$, h is rewired to j . With the remaining probability β , if $w_i = 0$, a link is formed between h and j , the link between i and h is maintained, and $w'_i = 1$.
- 2) $r = r_2$: If $|R_i|, |R_j| \geq 2$, both i and j choose one neighbor, $h \neq j \in R_i$ and $f \neq i \in R_j$, uniformly at random. If neither h nor f is linked to both i and j , then r_2 is executed by rewiring h to j and f to i .
- 3) $r = r_3$: If $d_i > d_j$, $|R_i| \geq 2$, and $w_i = 1$, then i chooses a neighbor, $h \neq j \in R_i$, uniformly at random. If $h \in \mathcal{N}_j$, then the link between i and h is removed and $w'_i = 0$.
- 4) $r = r_4$: i and j swap their binary flags, i.e. $w'_i = w_j$ and $w'_j = w_i$.

A feasible iteration of the algorithm on a network is illustrated in Fig. 3, where $\mathcal{G} = (V, E, w)$ is transformed into $\mathcal{G}' = (V, E', w')$ in one time step.

Algorithm II: Distributed Implementation of Φ^*

```

1 : initialize:  $\mathcal{G} = (V, E, w = \mathbf{0})$  connected,  $\epsilon, \beta \in (0, 1)$ 
2 : repeat
3 :   Each agent,  $i$ , is active with probability  $1 - \epsilon$ .
4 :   Each active  $i$  picks a random  $j \in \mathcal{N}_i$ .
5 :   For each  $i \in V$ ,  $R_i = \{i' \in \mathcal{N}_i \mid i' \text{ picked } i\}$ .
6 :   for (each  $(i, j)$  s.t.  $i \in R_j, j \in R_i, d_i \geq d_j$ )
7 :      $\max\{i, j\}$  picks a random  $r \in \Phi^*$ .
8 :     if ( $r = r_1, d_i > d_j, |R_i| \geq 2$ )
9 :        $i$  picks a random  $h \in R_i \setminus \{j\}$ .
10 :      if  $((j, h) \notin E)$ 
11 :         $i$  picks a random  $\beta' \in [0, 1]$ .
12 :        if ( $\beta' \geq \beta$ )
13 :           $E = (E \setminus \{(i, h)\}) \cup \{(j, h)\}$ .
14 :        else if ( $w_i = 0$ )
15 :           $E = E \cup \{(j, h)\}, w_i = 1$ .
16 :        end if
17 :      end if
18 :     else if ( $r = r_2, |R_i| \geq 2, |R_j| \geq 2$ )
19 :        $i$  picks a random  $h \in R_i \setminus \{j\}$ .
20 :        $j$  picks a random  $f \in R_j \setminus \{i\}$ .
21 :       if  $((i, f) \notin E, (j, h) \notin E)$ 
22 :          $E = (E \setminus \{(i, h), (j, f)\}) \cup \{(i, f), (j, h)\}$ .
23 :       else if ( $r = r_3, d_i > d_j, |R_i| \geq 2, w_i = 1$ )
24 :          $i$  picks a random  $h \in R_i \setminus \{j\}$ .
25 :         if  $((j, h) \in E)$ 
26 :            $E = E \setminus \{(i, h)\}, w_i = 0$ .
27 :         end if
28 :       else if ( $r = r_4$ )
29 :         Swap  $w_i$  and  $w_j$ .
30 :       end if
31 :     end for
32 : end repeat

```




Fig. 3. Some steps of a feasible iteration of Algorithm II on \mathcal{G} in (a) resulting in \mathcal{G}' in (c) along with the probabilities of the corresponding events. In this example, each node other than 8 is active and picks a neighbor as illustrated in (b), where each arrow is pointed from a node to its chosen neighbor. Accordingly, (1,3) and (4,6) are the matched pairs. With a joint probability of $1/16$, node 3 picks r_1 , and node 6 picks r_2 as the candidate rules for their respective matchings. Furthermore, since $R_3 \setminus \{1\} = \{5\}$, $R_4 \setminus \{6\} = \{2\}$, and $R_6 \setminus \{4\} = \{7\}$; node 3 picks node 5; node 4 picks node 2; and node 6 picks node 7 to rewire. Finally, the edge between node 3 and 5 is maintained with a probability β ($w_3 = 0$). Hence, given (b), \mathcal{G}' can emerge with a probability of $\beta/16$.

Note that Algorithm II maintains connectivity due to Lemma 4.3. Let $k = \bar{d}(\mathcal{G}(0))$ be the average degree of the initial graph, $\mathcal{G}(0) = (V, E(0), \mathbf{0})$. Then, the average degree remains in $[k, k+2]$ due to Corollary 4.5. Furthermore, Algorithm II is memoryless since each iteration only depends on the current graph, and the probability of any feasible transition is independent of the past. Hence, it induces a Markov chain over the state space, $\mathbb{G}_{n,[k,k+2]}$, consisting of connected labeled graphs with n nodes and average degrees contained in $[k, k+2]$, i.e.

$$\mathbb{G}_{n,[k,k+2]} = \{\mathcal{G} = (V, E, w) \mid |V| = n, |E| = 0.5kn + \mathbf{1}^T w\}, \quad (19)$$

where $w_i \in \{0, 1\}$ for every $i \in V$.

Let $\mu(t)$ denote the probability distribution over the state space at time t . Then, $\mu(t)$ satisfies

$$\mu^T(t+1) = \mu^T(t)P^*, \quad (20)$$

where P^* is the corresponding probability transition matrix. Accordingly, the probability of transition from any \mathcal{G} to any \mathcal{G}' is denoted by $P^*(\mathcal{G}, \mathcal{G}')$.

By following Algorithm II, the agents concurrently modify

their neighborhoods in accordance with Φ^* such that any feasible transformation occurs with a non-zero probability. In other words, for any pair of graphs, $\mathcal{G}, \mathcal{G}' \in \mathbb{G}_{n,[k,k+2]}$, the corresponding transition probability, $P^*(\mathcal{G}, \mathcal{G}')$, is non-zero if it is possible to transform \mathcal{G} into \mathcal{G}' in a single time step via Φ^* .

Lemma 5.1 *Let $\mathcal{G}, \mathcal{G}' \in \mathbb{G}_{n,[k,k+2]}$ be any pair of graphs. If \mathcal{G}' can be reached from \mathcal{G} in one step via Φ^* , then $P^*(\mathcal{G}, \mathcal{G}') > 0$.*

Proof:

Let \mathcal{G}' be reachable from \mathcal{G} in one step via Φ^* , and let $G = \{g_1, g_2, \dots\}$ denote the corresponding set of disjoint subgraphs of \mathcal{G} to be transformed to reach \mathcal{G}' . Note that for each $g \in G$ there is a $r = (g_l, g_r) \in \Phi^*$ satisfying $g \simeq g_l$. In order to show that $P^*(\mathcal{G}, \mathcal{G}') > 0$, we present a feasible flow of Algorithm II that transforms \mathcal{G} by applying the corresponding $r \in \Phi^*$ to each $g \in G$. For each $g \in G$, let each node in g be active and pick a neighbor as illustrated in Fig. 4.

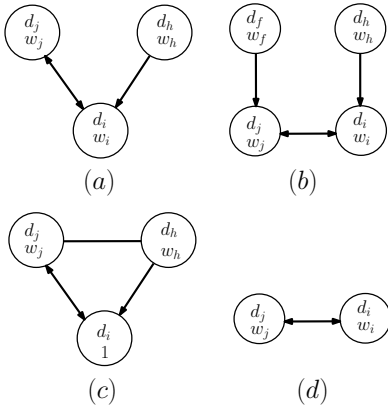


Fig. 4. An arrow is pointed from each agent to the neighbor it picked. For each $g \in G$, the nodes in g have non-zero probability to pick their neighbors as shown in (a) if $r = r_1$, (b) if $r = r_2$, (c) if $r = r_3$, and (d) if $r = r_4$.

Furthermore, let any node that is not included in any $g \in G$ be inactive, which ensures that only the subgraphs in G will be transformed. Finally, let each g pick the corresponding r as the candidate ruler to execute. In that case, the agents are guaranteed to only apply the corresponding $r \in \Phi^*$ to each $g \in G$. Hence, the corresponding transformation has a non-zero probability in P^* . ■

The state space, $\mathbb{G}_{n,[k,k+2]}$ can be represented as the union of two disjoint sets, $\mathbb{G}_{n,[k,k+2]}^0$ (regular graphs) and $\mathbb{G}_{n,[k,k+2]}^+$ (non-regular graphs), defined as

$$\mathbb{G}_{n,[k,k+2]}^0 = \{\mathcal{G} \in \mathbb{G}_{n,[k,k+2]} \mid f(\mathcal{G}) = 0\}, \quad (21)$$

$$\mathbb{G}_{n,[k,k+2]}^+ = \mathbb{G}_{n,[k,k+2]} \setminus \mathbb{G}_{n,[k,k+2]}^0. \quad (22)$$

In light of Lemma 4.10, $\mathbb{G}_{n,[k,k+2]}^0 \neq \emptyset$ for $k > 2$. Let M

denote the set of all such m , i.e.

$$M = \{m \in \mathbb{N} \mid k \leq m \leq k+2, 0.5m|V| \in \mathbb{N}\}. \quad (23)$$

The set of regular graphs, $\mathbb{G}_{n,[k,k+2]}^0$, can be written as the union of some disjoint sets as

$$\mathbb{G}_{n,[k,k+2]}^0 = \bigcup_{m \in M} \mathbb{G}_{n,m}^0, \quad (24)$$

where each $\mathbb{G}_{n,m}^0$ is the set of m -regular graphs defined as

$$\mathbb{G}_{n,m}^0 = \{\mathcal{G} \in \mathbb{G}_{n,[k,k+2]}^0 \mid \bar{d}(\mathcal{G}) = m\}. \quad (25)$$

In the remainder of this section, the limiting behavior of P^* is analyzed. In light of Lemma 4.10, if $k > 2$, then any connected $\mathcal{G}(0) = (V, E(0), \mathbf{0})$ almost surely converges to $\mathbb{G}_{n,[k,k+2]}^0$. Next, we show that, for each $m \in M$, $\mathbb{G}_{n,m}^0$ is a closed communicating class.

Lemma 5.2 *If $k > 2$, then $\mathbb{G}_{n,m}^0$ is a closed communicating class of P^* for each $m \in M$.*

Proof: In Light of Lemma 4.11, only r_2 and r_4 is applicable on any m -regular graph, and any resulting graph is m -regular. Hence, once the system reaches any $\mathbb{G}_{n,m}^0$, it stays in $\mathbb{G}_{n,m}^0$. Furthermore, in light of Theorem 4.2, any m -regular structure can be reached from any other via r_2 . Moreover, since all the graphs in $\mathbb{G}_{n,m}^0$ are connected, any permutation of the elements in w is also reachable via r_4 for each graph structure in $\mathbb{G}_{n,m}^0$. Hence, each $\mathbb{G}_{n,m}^0$ is a closed communicating class of P^* . ■

The limiting behavior of Algorithm II depends on the limiting behavior of P^* in each closed communicating class. Note that for each $\mathbb{G}_{n,m}^0$, there is a unique stationary distribution, μ_m^* , of P^* whose support is $\mathbb{G}_{n,m}^0$. Next, it is shown that each μ_m^* is a limiting distribution that is uniform over $\mathbb{G}_{n,m}^0$. To this end, first it is shown that the transitions between any two regular graphs are symmetric.

Lemma 5.3 *For any $m \in M$, let $\mathcal{G}, \mathcal{G}' \in \mathbb{G}_{n,m}^0$ be any pair of m -regular graphs. Then,*

$$P^*(\mathcal{G}, \mathcal{G}') = P^*(\mathcal{G}', \mathcal{G}). \quad (26)$$

Proof: For any $m \in M$, let $\mathcal{G}, \mathcal{G}' \in \mathbb{G}_{n,m}^0$ be any pair of m -regular graphs. Note that r_1 and r_3 are not applicable on any graph in $\mathbb{G}_{n,m}^0$ since all the nodes have equal degrees. Hence, any transition from \mathcal{G} to \mathcal{G}' is only via r_2 or r_4 . Since both r_2 and r_4 are reversible, $P^*(\mathcal{G}, \mathcal{G}') > 0$ if and only if $P^*(\mathcal{G}', \mathcal{G}) > 0$.

Let us consider an arbitrary execution of Algorithm II, where \mathcal{G} is transformed into \mathcal{G}' . Let u be the corresponding vector of randomly picked neighbors in line 4 of Algorithm II (let $u_i = \text{null}$ if i is inactive). For each node, i , let $R_i(u)$ be the set of nodes that picked i , and let $M(u) = \{(i, j) \mid i \in R_j(u), j \in R_i(u)\}$ be the set of matched pairs. In the remainder of the proof, it is shown that for any such feasible execution there exists an equally likely execution of Algorithm II that transforms \mathcal{G}' back to \mathcal{G} . For $\mathcal{G}' = (V, E', w')$, consider the vector, u' , whose entries are

$$u'_i = \begin{cases} u_i & \text{if } u_i = \text{null or } (i, u_i) \in E', \\ u_{u_i} & \text{otherwise.} \end{cases} \quad (27)$$

Note that $\Pr[u] = \Pr[u']$ since the inactive nodes will remain inactive with the same probability, and each active node picks a neighbor uniformly at random. Furthermore, $M(u) = M(u')$. Let each $(i, j) \in M(u)$ randomly choose the same candidate rule, r , they executed in the transition from \mathcal{G} to \mathcal{G}' . As such, if $r = r_4$, then i and j will swap w_i and w_j back. On the other hand, if $r = r_2$, then i and j will reverse the neighbor-swapping in the transition from \mathcal{G} to \mathcal{G}' with the same probability (lines 12-17 in Algorithm II) since $|R_i(u)| = |R_i(u')|$ and $|R_j(u)| = |R_j(u')|$ for every $(i, j) \in M(u)$. As such, all the local transformations from \mathcal{G} to \mathcal{G}' can be reversed with the same probability under Algorithm II. Consequently, $P^*(\mathcal{G}, \mathcal{G}') = P^*(\mathcal{G}', \mathcal{G})$. ■

Lemma 5.4 *For any $m \in M$, there is a unique limiting distribution, μ_m^* , of P^* satisfying*

$$\mu_m^*(\mathcal{G}) = \begin{cases} 1/|\mathbb{G}_{n,m}^0| & \text{if } \mathcal{G} \in \mathbb{G}_{n,m}^0, \\ 0 & \text{otherwise.} \end{cases} \quad (28)$$

Proof: For any $m \in M$, $\mathbb{G}_{n,m}^0$ is a closed communicating class due to Lemma 5.2. As such, for each $\mathbb{G}_{n,m}^0$, there exists a unique stationary distribution, μ_m^* , whose support is $\mathbb{G}_{n,m}^0$. Let P_m^* be the $|\mathbb{G}_{n,m}^0|$ by $|\mathbb{G}_{n,m}^0|$ probability transition matrix that only represents the transitions within $\mathbb{G}_{n,m}^0$. Due to Lemma 5.2, P_m^* is irreducible. Also P_m^* is aperiodic since $P^*(\mathcal{G}, \mathcal{G}) > 0$ for every \mathcal{G} (for instance, there is a non-zero probability of all the nodes being inactive). As such, the corresponding stationary distribution, μ_m^* , is a limiting distribution. Furthermore, due to Lemma 5.3, P_m^* is symmetric, and it is consequently doubly stochastic. As a result, μ_m^* is uniform over $\mathbb{G}_{n,m}^0$. ■

Theorem 5.5 *For any connected $\mathcal{G} = (V, E, \mathbf{0})$ satisfying $\bar{d}(\mathcal{G}) > 2$, P^* leads to a limiting distribution, $\mu_{\mathcal{G}}^*$, given as*

$$\mu_{\mathcal{G}}^* = \sum_{m \in M} \Pr[\mathcal{G} \rightarrow \mathbb{G}_{n,m}^0] \mu_m^*, \quad (29)$$

where $\Pr[\mathcal{G} \rightarrow \mathbb{G}_{n,m}^0]$ is the probability that the chain initialized at \mathcal{G} ever enters the set of m -regular graphs, $\mathbb{G}_{n,m}^0$, and each μ_m^* is uniform over its support, $\mathbb{G}_{n,m}^0$.

Proof: Due to Lemma 4.10, the Markov chain initialized at any connected $\mathcal{G} = (V, E, \mathbf{0})$ satisfying $\bar{d}(\mathcal{G}) > 2$ almost surely enters a closed communicating class of m -regular graphs, $\mathbb{G}_{n,m}^0$, such that $m \in M$. Furthermore, each $\mathbb{G}_{n,m}^0$ has a uniform limiting distribution, μ_m^* as given in Lemma 5.4. Hence, P^* leads to the convex combination of limiting distributions, μ_m^* , where each μ_m^* is weighted by the probability that the chain starting at \mathcal{G} ever enters $\mathbb{G}_{n,m}^0$. ■

In light of Theorem 5.5, Algorithm II asymptotically transforms any connected initial graph, $\mathcal{G} = (V, E, \mathbf{0})$, satisfying

$\bar{d}(\mathcal{G}) > 2$ into a random m -regular graph for some $m \in M$ as given in (23). When M is not a singleton, the probability of reaching each $\mathbb{G}_{n,m}^0$ depends on the initial graph as well as the algorithm parameters $\epsilon, \beta \in (0, 1)$. However, the graphs observed in the limit are guaranteed to be expanders since $\bar{d}(\mathcal{G}) > 2$ implies $m \geq 3$ for every $m \in M$.

VI. SIMULATION RESULTS

In this section, we present some simulation results for the proposed scheme. An arbitrary connected graph, $\mathcal{G}(0) = (V, E, \mathbf{0})$, is generated with 100 nodes and 135 edges. As such, the initial average degree is $\bar{d}(\mathcal{G}(0)) = 2.7$. The interaction graph is evolved via Algorithm II with $\epsilon = \beta = 0.01$ for a period of 100000 steps. Since $\bar{d}(\mathcal{G}(0)) = 2.7$, the average degree is guaranteed to remain in $[2.7, 4.7]$ for any feasible trajectory, as given in Corollary 4.5. Since the number of nodes is even, both integers in this interval are feasible values for the average degree, i.e. $M = \{3, 4\}$. In light of Theorem 5.5, the graph is expected to become either a random 3-regular graph or a random 4-regular graph. In the presented simulation, the graph reaches a 3-regular graph after 45123 time steps. Note that once the system enters $\mathbb{G}_{100,3}^0$, both $f(\mathcal{G}(t))$ and $\bar{d}(\mathcal{G}(t))$ are stationary. Fig. 5 illustrates $\mathcal{G}(0)$ and $\mathcal{G}(100000)$, which have the algebraic connectivities of 0.02 and 0.25, respectively. The values of the degree range, $f(\mathcal{G}(t))$, and the average degree, $\bar{d}(\mathcal{G}(t))$, for the first 50000 steps are illustrated in Fig. 6. Once a 3-regular graph is reached, the graph keeps randomizing in $\mathbb{G}_{100,3}^0$ via r_2 and r_4 . The evolution of the algebraic connectivity throughout this simulation is shown in Fig. 7. As expected from random 3-regular graphs, the algebraic connectivity of the network is observed to be at least $3 - 2\sqrt{2}$ with a very high probability after a sufficient amount of time.

VII. CONCLUSIONS

In this paper, we presented a decentralized graph reconfiguration scheme for building robust multi-agent networks. In particular, we focused on the connectivity properties of the interaction graph as the robustness measure. Accordingly, we provided a decentralized method for transforming interaction graphs into well-connected graphs with a similar sparsity as the initial graph. More specifically, the proposed solution produces random m -regular graphs in the limit. Such graphs are expanders for any $m \geq 3$, i.e. they have the algebraic connectivity and expansion ratios bounded away from zero regardless of the network size.

The proposed method was incrementally built in the paper. First, a single-rule grammar, Φ_R , was designed for balancing the degree distribution in any connected network with an initial average degree k . Φ_R transforms the initial graph into a connected k -regular graph if $k \in \mathbb{N}$. If $k \notin \mathbb{N}$, then the degree range converges to 1 via Φ_R . Next, Φ_R was extended to ensure that the interaction graph does not converge to a poorly-connected k -regular graph. To this end, a local link randomization rule was added to Φ_R . The resulting grammar, Φ_{RR} , transforms the initial graph into a connected random k -regular graph if $k \in \mathbb{N}$. Finally, Φ_{RR} was extended to

obtain random regular graphs in the most general case, i.e. even when $k \notin \mathbb{N}$. For any $k > 2$, the resulting grammar, Φ^* , leads to a connected random m -regular graph such that $m \in [k, k+2]$. Note that $k > 2$ implies $m \geq 3$. Hence, the corresponding random m -regular graphs are expanders with a similar sparsity as the initial graph. Some simulation results were also presented in the paper to demonstrate the performance of the proposed method.

REFERENCES

- [1] A. H. Dekker and B. D. Colbert, "Network robustness and graph topology," in *Australasian Conference on Computer Science*, pp. 359–368, 2004.
- [2] A. Jamakovic and S. Uhlig, "On the relationship between the algebraic connectivity and graph's robustness to node and link failures," in *EuroNGI Conference on Next Generation Internet Networks*, pp. 96–102, 2007.
- [3] R. Olfati-Saber, "Ultrafast consensus in small-world networks," in *American Control Conference*, pp. 2371–2378, 2005.
- [4] A. Rahmani, M. Ji, M. Mesbahi, and M. Egerstedt, "Controllability of multi-agent systems from a graph-theoretic perspective," *SIAM Journal on Control and Optimization*, vol. 48, no. 1, pp. 162–186, 2009.
- [5] A. Y. Yazıcıoğlu, W. Abbas, and M. Egerstedt, "A tight lower bound on the controllability of networks with multiple leaders," in *IEEE Conference on Decision and Control*, pp. 1978–1983, 2012.
- [6] M. Mesbahi and M. Egerstedt, *Graph theoretic methods in multiagent networks*. Princeton University Press, 2010.
- [7] A. Barrat, M. Barthélemy, and A. Vespignani, *Dynamical processes on complex networks*. Cambridge University Press Cambridge, 2008.
- [8] G. F. Young, L. Scardovi, and N. E. Leonard, "Robustness of noisy consensus dynamics with directed communication," in *American Control Conference (ACC)*, pp. 6312–6317, 2010.
- [9] W. Abbas and M. Egerstedt, "Robust graph topologies for networked systems," in *IFAC Workshop on Distributed Estimation and Control in Networked Systems*, pp. 85–90, 2012.
- [10] M. Fiedler, "Algebraic connectivity of graphs," *Czechoslovak Mathematical Journal*, vol. 23, no. 2, pp. 298–305, 1973.
- [11] D. J. Klein and M. Randić, "Resistance distance," *Journal of Mathematical Chemistry*, vol. 12, no. 1, pp. 81–95, 1993.
- [12] M. S. Pinsker, "On the complexity of a concentrator," in *7th International Teletraffic Conference*, vol. 4, pp. 318/1–318/4, 1973.
- [13] N. Alon, "Eigenvalues and expanders," *Combinatorica*, vol. 6, no. 2, pp. 83–96, 1986.
- [14] S. Hoory, N. Linial, and A. Wigderson, "Expander graphs and their applications," *Bulletin of the American Mathematical Society*, vol. 43, no. 4, pp. 439–561, 2006.
- [15] L. Blume, D. Easley, J. Kleinberg, R. Kleinberg, and É. Tardos, "Which networks are least susceptible to cascading failures?," in *IEEE Annual Symposium on Foundations of Computer Science*, pp. 393–402, 2011.
- [16] A. Ganesh, L. Massoulié, and D. Towsley, "The effect of network topology on the spread of epidemics," in *IEEE INFOCOM*, vol. 2, pp. 1455–1466, 2005.
- [17] J. Friedman, "A proof of alon's second eigenvalue conjecture," in *ACM Symposium on Theory of Computing*, pp. 720–724, 2003.
- [18] A. Y. Yazıcıoğlu, M. Egerstedt, and J. S. Shamma, "Decentralized degree regularization for multi-agent networks," in *IEEE Conference on Decision and Control*, pp. 7498–7503, 2013.
- [19] A. Y. Yazıcıoğlu, M. Egerstedt, and J. S. Shamma, "Decentralized formation of random regular graphs for robust multi-agent networks," in *IEEE Conference on Decision and Control*, pp. 595–600, 2014.
- [20] S. P. Borgatti, "Centrality and network flow," *Social Networks*, vol. 27, no. 1, pp. 55–71, 2005.
- [21] D. Koschützki, K. A. Lehmann, L. Peeters, S. Richter, D. Tenfelde-Podehl, and O. Zlotowski, "Centrality indices," in *Network Analysis*, pp. 16–61, Springer, 2005.
- [22] D. Acemoglu, V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi, "The network origins of aggregate fluctuations," *Econometrica*, vol. 80, no. 5, pp. 1977–2016, 2012.
- [23] H. Jeong, S. P. Mason, A.-L. Barabási, and Z. N. Oltvai, "Lethality and centrality in protein networks," *Nature*, vol. 411, no. 6833, pp. 41–42, 2001.
- [24] P. Holme, B. J. Kim, C. N. Yoon, and S. K. Han, "Attack vulnerability of complex networks," *Physical Review E*, vol. 65, no. 5, p. 056109, 2002.
- [25] Z. Wang, A. Scaglione, and R. J. Thomas, "Electrical centrality measures for electric power grid vulnerability analysis," in *IEEE Conference on Decision and Control*, pp. 5792–5797, 2010.
- [26] C. M. Schneider, A. A. Moreira, J. S. Andrade, S. Havlin, and H. J. Herrmann, "Mitigation of malicious attacks on networks," *Proceedings of the National Academy of Sciences*, vol. 108, no. 10, pp. 3838–3841, 2011.
- [27] A. Beygelzimer, G. Grinstein, R. Linsker, and I. Rish, "Improving network robustness by edge modification," *Physica A: Statistical Mechanics and its Applications*, vol. 357, no. 3, pp. 593–612, 2005.
- [28] O. Reingold, S. Vadhan, and A. Wigderson, "Entropy waves, the zig-zag graph product, and new constant-degree expanders," *Annals of Mathematics*, pp. 157–187, 2002.
- [29] M. Capalbo, O. Reingold, S. Vadhan, and A. Wigderson, "Randomness conductors and constant-degree lossless expanders," in *ACM Symposium on Theory of Computing*, pp. 659–668, 2002.
- [30] E. Rozenman and S. Vadhan, "Derandomized squaring of graphs," in *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*, pp. 436–447, Springer, 2005.
- [31] M. Morgenstern, "Existence and explicit constructions of $q+1$ regular ramanujan graphs for every prime power q ," *Journal of Combinatorial Theory, Series B*, vol. 62, no. 1, pp. 44–62, 1994.
- [32] R. Olfati-Saber, "Algebraic connectivity ratio of ramanujan graphs," in *American Control Conference*, pp. 4619–4624, 2007.
- [33] N. C. Wormald, "Models of random regular graphs," *London Mathematical Society Lecture Note Series*, pp. 239–298, 1999.
- [34] B. Mohar, "Isoperimetric numbers of graphs," *Journal of Combinatorial Theory, Series B*, vol. 47, no. 3, pp. 274–291, 1989.
- [35] B. Bollobás, "A probabilistic proof of an asymptotic formula for the number of labelled regular graphs," *European Journal of Combinatorics*, vol. 1, no. 4, pp. 311–316, 1980.
- [36] A. Steger and N. C. Wormald, "Generating random regular graphs quickly," *Combinatorics Probability and Computing*, vol. 8, no. 4, pp. 377–396, 1999.
- [37] C. Law and K.-Y. Siu, "Distributed construction of random expander networks," in *IEEE International Conference on Computer Communications*, pp. 2133–2143, 2003.
- [38] M. Jerrum and A. Sinclair, "Fast uniform generation of regular graphs," *Theoretical Computer Science*, vol. 73, no. 1, pp. 91–100, 1990.
- [39] P. Mahdian and C. Schindelhauer, "Peer-to-peer networks based on random transformations of connected regular undirected graphs," in *ACM Symposium on Parallelism in Algorithms and Architectures*, pp. 155–164, 2005.
- [40] E. Klavins, R. Ghrist, and D. Lipsky, "A grammatical approach to self-organizing robotic systems," *IEEE Transactions on Automatic Control*, vol. 51, no. 6, pp. 949–962, 2006.
- [41] A. Kashyap, T. Başar, and R. Srikant, "Quantized consensus," *Automatica*, vol. 43, no. 7, pp. 1192–1203, 2007.
- [42] A. Nedic, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis, "On distributed averaging algorithms and quantization effects," *IEEE Transactions on Automatic Control*, vol. 54, no. 11, pp. 2506–2517, 2009.




Fig. 5. Agents transform the initial interaction graph in (a), which has an average degree of 2.7, into a random 3-regular graph such as the one in (b) by following Algorithm II.




Fig. 6. The average degree, $\bar{d}(\mathcal{G}(t))$, and the degree range, $f(\mathcal{G}(t))$, for the first 50000 time steps of the simulation. Once a regular graph (i.e., $f(\mathcal{G}(t)) = 0$) is reached, both $\bar{d}(\mathcal{G}(t))$ and $f(\mathcal{G}(t))$ remain stationary under Algorithm II.




Fig. 7. The algebraic connectivity, $\alpha(\mathcal{G}(t))$, as the initial graph in Fig. 5a evolves via Algorithm II. After a sufficiently large amount of time, $\alpha(\mathcal{G}(t))$ rarely drops below $3 - 2\sqrt{2}$ (marked with a horizontal solid line), as expected from random 3-regular graphs.